

ABSTRACT

This invention provides to large-scale operating systems supporting multiple classes of tasks a method to allocate processor resources to the classes according to a probability model. It is useful in preventing a class from being denied resources by giving it positive measure in the model. A first stage probability algorithm assigns classes of tasks to an IP resource available to a given scheduler queue. Each class is allocated a probability level in a lottery-type draw in this first stage. In preferred embodiments, a second stage probability algorithm is used to assign tasks within a class to an available processor resource. This second stage algorithm is biased in a feedback manner by task resource usage history. Tasks of extreme high or low priority may avoid the probabilistic mechanisms in preferred embodiments.